

REMARKS

Claims 1 stands objected to and correction is requested. Claims 1-4, 6, 10-14 and 15-18 stand rejected under 35 USC §103(a) as being unpatentable over Kelman, U.S. patent 6,671,820 in view of Stai et al., U.S. patent 6,401,128 and further in view of Haren, U.S. patent 6,557,060, and further in view of newly cited Fredericks et al., U.S. patent 6,347,334. Claims 8, and 9 stand rejected under 35 USC §103(a) as being unpatentable over Kelman, U.S. patent 6,671,820 in view of Stai et al., U.S. patent 6,401,128 and Haren, U.S. patent 6,557,060 and Fredericks et al., U.S. patent 6,347,334, and further in view of Panas et al., U.S. patent 6,473,857.

Claims 1, and 11 have been amended to accommodate the Examiner's noted objection and to more clearly define the invention.

Reconsideration and allowance of each of the claims 1-4, 6 and 8-18, as amended, is respectfully requested. Applicants respectfully submit that each of the pending claims 1-4, 6 and 8-18, as amended, is patentable over the references of record.

Summary of Scope and Content of the Prior Art

Kelman, U.S. patent 6,671,820 discloses a system and method for preventing the corruption of networked storage devices during the process of transferring backup data to a server on the network that has suffered data loss as the result of a disaster. Column 3, line 51 - column 4, line 2 states: The system and method described herein include a LUN masking driver. The LUN masking driver is preferably contained on an emergency diskette that is to be used during the recovery

process for loading vital device drivers onto the affected server so that the affected server may boot and connect to the SAN. During the recovery process, the LUN masking driver will load when the operating system boots up, after the SAN HBA driver loads and before the normal file systems load. The LUN masking driver scans all devices visible on the SAN and uses SCSI inquiry commands to determine which devices are dedicated backup storage devices. The LUN masking driver then masks all devices that are not dedicated backup storage devices. Thus, only dedicated backup storage devices are visible to software that boots up after the LUN masking driver completes its function. Consequently, the operating system's file systems never see the storage devices that are not dedicated backup storage devices. As a result, the affected server cannot access the storage devices and cause data corruption. Column 6, line 63 - column 7, line 36 states:

After the HBA and disk drivers have been loaded, the LUN masking driver loads onto the affected server at step 42. The LUN masking driver then issues a SCSI protocol inquiry of the devices on shared storage network 10. The SCSI protocol inquiry is intended to determine two things: first, whether the interrogated device is a SCSI device, and second, what type of SCSI device. When a SCSI type device receives a SCSI protocol inquiry, the SCSI device returns a signal to the device or software that issued the inquiry. This signal returned by the SCSI device includes peripheral-type information that identifies the device type. For example, if the SCSI device is a tape storage device, the peripheral-type information will identify the SCSI device as a tape storage device. At step 44, the LUN masking driver receives the peripheral-type information from the SCSI devices on shared storage network 10. For example, if the LUN masking driver issues a SCSI protocol inquiry of the devices in shared storage network 10, then storage devices 24, dedicated backup storage devices 28, and any other SCSI device comprising computer network 14 will return a signal to the LUN masking driver that contains peripheral-type information. The peripheral-type information sent by storage device 24 will include information identifying the storage device 24 as a hard disk drive, for example, and may include specific information such as the manufacturer, make, or model of storage device 24. The dedicated backup storage devices 28 will return similar information to identify its device type, i.e., a tape drive. Once the LUN masking driver has received the peripheral-type information, the

LUN masking driver will be able to identify what SCSI devices on shared storage network 10 are dedicated storage devices 28.

At step 46, the operating system loads. When the operating system initially loads onto the affected server, the operating system will communicate with the disk driver to identify the storage devices 24 that are located on the shared storage network 10. Accordingly, at step 48a, the operating system issues a command to identify all of the available LUNs on the storage network 10. Unless the LUNs are masked, the disk driver will respond with all of the LUN addresses which will, as discussed above, create the risk of data corruption.

Stai et al., U.S. patent 6,401,128 discloses a system and method for sending frames between a public device and a private device comprise a phantom device mapping, an address translation, a frame payload translation, and a CRC regeneration. The system and method assign a phantom AL_PA for the public device and establishes a phantom device mapping between the phantom AL_PA and the public device's Port_ID. With the phantom device mapping, the disclosed system directs all communication between the public device and the private loop device as if the communication were between a phantom device and the private device. Specifically, the system and method comprise a public-to-private address translation in one direction and a private-to-public address translation in the other direction. During the public-to-private address translation process, the source address of the frame is converted to a phantom AL_PA. The public-to-private address translation uses a Port_ID to phantom AL_PA mapping table and finds an entry where the Port_ID matches the source address. The public-to-private address translation replaces the source address of the frame with the phantom AL_PA of the matched entry, and the destination address with the AL_PA only of the destination device. The private-to-

public address translation replaces the destination address of the frame with the Port_ID of the matched entry, and the source address with the fabric assigned address of the private device. As set forth at column 6, starting at line 25: Every frame between public device 106A and private loop device 110E requires an address translation, either a public-to-private or a private-to-public address translation. In addition, if the frame content (called payload) contains any address information, it is also changed accordingly. Typically, these types of frames are Extended Link Services (ELS) in Fibre Channel, and they are either ELS request or response frames. FIG. 4 is a functional block diagram of a preferred embodiment of an ELS request payload translation. An ELS request payload translation process is required if the source address and/or the destination address of a frame being transmitted between a public device and a private device are part of the payload. An ELS request payload translation may be performed during a public-to-private translation or during a private-to-public translation. During the ELS request payload translation process, the frame type and Extended Link Services command code are examined to determine if payload translation is required. These are examined using an ELS request payload table which stores information on the frame type, command codes, and the information of fields to be modified. If payload translation is required, one or more fields in the frame payload is translated according to the ELS request payload table. If an ELS request payload translation is performed, the information for that frame is stored in a request payload cross-reference table. The request payload cross-reference table can then be used during an ELS response payload translation as described with reference to FIG. 5.

Haren, U.S. patent 6,557,060 discloses a host expansion bridge where data is converted from a first granularity to a second granularity different from the first granularity. The ratio "n" of the second granularity of the data to the first granularity of the data is determined as a power of 2. The least significant n bits of the beginning alignment of the data are added to the least significant n bits of the beginning count of the data, and the carry bit of the sum is designated as "c". A logical "OR" is performed of the bits of the resulting sum to obtain a value designated as "d". A number of data units, equal to the sum of "c" and "d", is added to the data. Column 4, lines 1-28 states:

A software stack may be provided in channel adapter 119 or 119' to access the network switching fabric 100 and information about fabric configuration, fabric topology and connection information. The operating system software (OS) of the processing system 110 may include a fabric bus driver and a fabric adapter device-specific driver utilized to establish communication with a remote fabric-attached agent (e.g., I/O controller) of another processing system connected to the network, and perform functions common to most drivers, including, for example, host-fabric adapter initialization and configuration, channel configuration, channel abstraction, resource management, fabric management service and operations, send/receive I/O transaction messages, remote direct memory access (rDMA) data transfers (e.g., read and write operations), queue management, memory registration, descriptor management, message flow control, and transient error handling and recovery. Such a software driver module may be written using high-level programming languages such as C, C++ and Visual Basic, and may be provided on a tangible medium, such as a memory device, magnetic disk (fixed, floppy, and removable), other magnetic media such as magnetic tapes; optical media such as CD-ROM disks, or via Internet download, which may be available for a network administrator to conveniently plug-in or download into an existing operating system (OS). Such a software driver module may also be bundled with the existing operating system which may be activated by a particular device driver.

Fredericks et al., U.S. patent 6,347,334 discloses a method for implementing a link level service in a computer network having a first port device and a second port device. Node identification data is stored in the second port device. A physical-layer communications coupling is provided between the first port device and

the second port device which may be a point-to-point, loop, or switched circuit connection. The first port device sends a request node identification (RNID) message addressed to the second port device. The second port device creates an accept message and copies stored node identification data into the accept message. The second port device sends the accept message to the first port device. Column 6, lines 9-44 state:

As described hereinbefore the RNID can optionally request node information from only nodes that support a specific FC-4 frame format (e.g., SCSI, SBCCS, and the like). In this way, the requesting device will not receive Accept message responses from devices that are not of interest. The ELS implementation preferably allows that if the recipient device does not support the requested node-identification data format specified in the RNID message, it should reply with a "link service reject" (LS_RJT) ELS message with a reason code of "unable to perform command request" and reason-code explanation of "unable to supply requested data".

In a preferred implementation, a RNID message is sent by an N_Port to its nearest neighbor node. The node attached to the other end of the link is hereinafter referred to as the "nearest neighbor node" and may be a node port (N_Port) of another device or a fabric port (F_Port). It is permissible, however, for a node to request node-identification data from any other node in the fabric. If the nearest neighbor node has been determined from the fabric logon information (FLOGI) to be a fabric node, then the RNID ELS message is sent to the Fabric Controller at the address hex `FFFFFD` as is well known. If the nearest neighbor node has been determined from the FLOGI to be an N_Port, then the RNID ELS is sent directly to that N_Port identified by its address identifier.

The fabric may acquire the node-identification data for all nodes attached to a switch's E_Ports, FL_Ports, or E_Ports with the address identifier of the attached node. Multiple switches in the same fabric may acquire the node-identification data of the node attached to the other end of their inter-switch link (ISL). For node-identification data acquisition between E_Port nodes, each switch 106 and 107 may issue and respond to RNID ELS messages across the ISL using Class F service (i.e. the ELS inherently becomes a switch fabric internal link services (SW_ILS)).

Panas et al., U.S. patent 6,473,857 discloses a method for centralized and managed loading of boot images into one or more processors that are part of a file

server for a mass storage system. In a computer system having at least one first controller, at least one input output processor (IOP), a first bus and a second bus, the present invention includes the steps of detecting readiness of the IOP to load a boot image, identifying across the first bus a location where the boot image will be loaded and loading the boot image across the second bus. The first controller may determine which of a plurality of boot images should be loaded. The first controller and the IOP may each have first and second processors, with communication between the first processors being across the first bus and boot images being accessed by the second processors across the second bus. On the IOP, the first processor may control power to the second processor and may monitor the status of the second processor, reporting across the first bus to the first controller's first processor regarding the status of the IOP's second processor. The boot image may be copied to memory local to the IOP's second processor or it may be made available across the second bus. The boot image supplied may be adapted to normal, diagnostic, crash dump or other purposes. The progress of IOP booting is tracked and monitored. As stated at column 4, lines 36-61: The connection options 130 are various methods of connecting servers and clients to the ISAN server 102A. The serial connections 140 support network management, modems for remote management, and uninterruptible power supply messages. The front panel connection 142 supports a management connection with the front panel display of the ISAN server 102A. The Ethernet connection 144 supports an Ethernet interface for management protocols and for data transfer. The network interface 146 is one of potentially many high speed interfaces on the server. In some embodiments, the

network interface 146 is a fibre channel interface with drivers for a fibre channel arbitrated loop (FC-AL). The network interface 146 may also include drivers for SCSI-3 over the fibre channel medium using fibre channel protocol (FCP). The hardware interface 126 provides interface specific hardware components. For example, the network interface 146 has a network interface specific set of software modules to support configuration, diagnostics, performance monitoring, and health and status monitoring. The operating system 124, the tables 116, and the interfaces 118-122 support the virtual device and storage routing functionality of the ISAN server 102A. These components of the ISAN server 102A route storage transactions among appropriate storage options 128 and the connection options 130 using configured sets of driver modules in the system.

Claims 1, and 11 are patentable

Applicants respectfully submit that each of the independent claims 1 and 11, as amended, is patentable over all the references of record including Kelman, Stai et al., Haren and Fredericks et al.

The present invention, as recited in independent claims 1 and 11, as amended, provides a novel storage area network (SAN) management and configuration method and apparatus via enabling in-band communications that solves a problem of some existing SAN arrangements. A problem exists in some known storage area network arrangements, for example, in a serial storage architecture (SSA), device driver writers and host based adapter (HBA) vendors provide a complex set of micro code calls. A management program would then interrogate the HBA, using micro code calls

specific to the particular HBA vendor and model, then interpret the results in a way that is specific to that particular HBA vendor and model. One problem with this arrangement is that an in-depth understanding is needed for every HBA model of every vendor, which in the case of Fibre Channel, is impractical. There are too many vendors and too many models to implement this approach. The present invention, as recited in independent claims 1 and 11, as amended, provides a pass through in said HBA device driver for passing communications to a device in the storage area network from said SAN management application, including at least one topology analysis command, and at least a transport pass through and an extended link service (ELS) pass through; each of said transport pass through and said extended link service (ELS) pass through being a binary pass through, each taking applied commands and passing said commands to said designated device in the storage area network as recited in amended independent claims 1 and 11.

Independent claim 1, as amended, recites a storage area network (SAN) management and configuration method via enabling in-band communications comprising the steps of: utilizing a SAN management application for managing and configuring the storage area network; said SAN management application communicates with at least one SAN-connected host system and communicates with a host bus adapter (HBA) device driver, and providing a pass through in said HBA device driver for passing communications to a designated device in the storage area network from said SAN management application including at least one topology analysis command; said at least one topology analysis command including a command to get interconnect

information and a command to get topology information; and providing said pass through includes providing at least a transport pass through and an extended link service (ELS) pass through; each of said transport pass through and said extended link service (ELS) pass through being a binary pass through, each taking applied commands and passing said commands to said designated device in the storage area network.

As presented, independent claim 1 further defines the transport pass through and the extended link service (ELS) pass through, reciting that each of said transport pass through and said extended link service (ELS) pass through being a binary pass through, each taking applied commands and passing said commands to said designated device in the storage area network.

Applicants respectfully submit that the references of record including Frederick et al. provide no suggestion nor motivation for providing a pass through including the transport pass through and the extended link service (ELS) pass through in said HBA device driver that are binary pass throughs, each taking applied commands and passing said commands to said designated device in the storage area network, as taught by Applicants and recited in independent claim 1, as amended. Contrary to the rejection, Frederick et al. provide no suggestion nor motivation of any binary pass throughs. Frederick et al. provide no suggestion nor motivation for any pass throughs taking applied commands and passing said commands to said designated device in the storage area network, as taught by Applicants and recited in independent claim 1, as amended. Thus, independent claim 1, as amended, is patentable.

Independent claim 11, as amended, recites a storage area network (SAN) management and configuration apparatus via enabling in-band communications comprising: a storage area network (SAN) management application for managing and configuring the storage area network via enabling in-band communications, said SAN management application for communicating with at least one SAN-connected host system; said SAN-connected host system including a management application agent for communicating with a host bus adapter (HBA) device driver; said HBA device driver for communicating with a device in the storage area network; said HBA device driver including at least one pass through service for passing a plurality of commands to said device in the storage area network; said commands including at least one topology analysis command; and said at least one topology analysis command including a command to get interconnect information and a command to get topology information; said pass through including at least a transport pass through and an extended link service (ELS) pass through; each of said transport pass through and said extended link service (ELS) pass through being a binary pass through, each taking applied commands and passing said commands to said designated device in the storage area network.

The references of record including Kelman, Stai et al., Haren, and Fredericks et al. considering the total teachings in combination, fail to suggest a storage area network (SAN) management and configuration apparatus via enabling in-band communications comprising a HBA device driver for communicating with a device in the storage area network; said commands including at least one topology analysis

command and said HBA device driver including at least one pass through service for passing a plurality of commands to said device in the storage area network; and each of said transport pass through and said extended link service (ELS) pass through being a binary pass through, each taking applied commands and passing said commands to said designated device in the storage area network. Thus, independent claim 11, as amended, is patentable.

Applicants respectfully submit that all the references of record including Kelman, Stai et al., Haren, Fredericks et al. and Panas et al., do not render obvious the claimed invention and each of the independent claims 1 and 11, as amended, is patentable. No objective teaching in the prior art or knowledge generally available to one of ordinary skill in the art suggests the claimed subject matter of independent claims 1 and 11, as amended. The references of record including Kelman, Stai et al., Haren, Fredericks et al. and Panas et al. fail to suggest or provide any objective teaching of a storage area network (SAN) management and configuration method and apparatus via enabling in-band communications as taught and claimed by Applicants in each of the independent claims 1 and 11, as amended.

Applicants respectfully submit that the teachings or suggestions found in the prior art including Kelman, Stai et al., Haren, Fredericks et al. and Panas et al. would not have been led one of ordinary skill in the art to the claimed invention. Applicants respectfully submit that each of the independent claims 1 and 11, as amended, is patentable.

Serial No. 09/657,234

Dependent claims 2-4, 6, 8-10, and 12-18 respectively depend from patentable independent claims 1 and 11 and further define the invention. Thus, each of the dependent claims 2-4, 6, 8-10, and 12-18 is likewise patentable.


Applicants have reviewed all the art of record, and respectfully submit that the claimed invention is patentable over all the art of record, including the references not relied upon by the Examiner for the rejection of the pending claims.

It is believed that the present application is now in condition for allowance and allowance of each of the pending claims 1-4, 6 and 8-18, as amended, is respectfully requested. Prompt and favorable reconsideration is respectfully requested.

If the Examiner upon considering this amendment should find that a telephone interview would be helpful in expediting allowance of the present application, the Examiner is respectfully urged to call the applicants' attorney at the number listed below.

Respectfully submitted,

By: _____


Joan Pennington
Reg. No. 30,885
Telephone: (312) 670-0736